arXiv:1204.1393v1 [cs.CV] 6 Apr 2012

Continuous Markov Random Fields for Robust Stereo Estimation

Koichiro Yamaguchi Tamir Hazan David McAllester Raquel Urtasun
Toyota Technological Institute at Chicago

March 2, 2013

Abstract

In this paper we present a novel slanted-plane MRF model which reasons jointly about occlusion boundaries as well as depth. We formulate the problem as the one of inference in a hybrid MRF composed of both continuous (i.e., slanted 3D planes) and discrete (i.e., occlusion boundaries) random variables. This allows us to define potentials encoding the ownership of the pixels that compose the boundary between segments, as well as potentials encoding which junctions are physically possible. Our approach outperforms the state-of-the-art on Middlebury high resolution imagery [1] as well as in the more challenging KITTI dataset [2], while being more efficient than existing slanted plane MRF-based methods, taking on average 2 minutes to perform inference on high resolution imagery.

1 Introduction

Over the past few decades we have witnessed a great improvement in performance of stereo algorithms. Most modern approaches frame the problem as inference on a Markov random field (MRF) and utilize global optimization techniques such as graph cuts or message passing [3] to reason jointly about the depth of each pixel in the image.

A leading approach to stereo vision uses slanted-plane MRF models which were introduced a decade ago in [4]. Most methods [5, 6, 7, 8, 9, 10] assume a fixed set of superpixels on a reference image, say the left image of the stereo pair, and model the surface under each superpixel as a slanted plane. The MRF typically has a robust data term scoring the assigned plane in terms of a matching score induced by the plane on the pixels contained in the superpixel. This data term often incorporates an explicit treatment of occlusion — pixels in one image that have no corresponding pixel in the other image [11, 12, 8, 13]. Slanted-plane models also typically include a robust smoothness term expressing the belief that the planes assigned to adjacent superpixels should be similar.

A major issue with slanted-plane stereo models is their computational complexity. For example, [13] reports an average of approximately one hour of computation for each low-resolution Middebury stereo pair. This makes these approaches impractical for applications such as robotics or autonomous driving. A main source of difficulty is that each plane is defined by three continuous parameters and inference for continuous MRFs with non-convex energies is computationally challenging.

This paper contains two contributions. First, we introduce the use of junction potentials, described below, into this class of models. Second, we show that particle methods can achieve strong performance with reasonable inference times on the high-resolution, in-the-wild KITTI dataset [2].

Junction potentials originate in early line labeling algorithms [14, 15]. These algorithms assign labels to the lines of a line drawing where the label indicate whether the line represents a discontinuity due to changes in depth (an occlusion), surface orientation (a corner), lighting (a shadow) or albedo (paint). A junction is a place where three lines meet. Only certain combinations of labels are physically realizable at junctions. The constraints on label combinations at junctions often force the labeling of the entire line drawing [14]. Here, as in recent work on monocular image interpretation [16, 17, 18], we label the boundaries between image segments rather than the lines of a line drawing with labels –"left occlusion", "right occlusion", "hinge" or

"coplanar". In our model the occlusion labels play a role in the data term where they are interpreted as expressing ownership of the pixels that compose the boundary between segments — an occlusion boundary is "owned" by the foreground object.

Our second contribution is to show that particle methods can be used to implement high performance inference in high resolution imagery with reasonable running time. Particle methods avoid premature commitment to any fixed quantization of continuous variables and hence allow a precise exploration of the continuous space. Our particle inference method is based on the recently developed particle convex belief propagation (PCBP) [19]. We learn the contribution of each potential via the primal-dual optimization framework of [20].

In the remainder of the paper we first review related work. We then introduce our continuous MRF model for stereo and show how to do learning and inference in this model. Finally, we demonstrate the effectiveness of our approach in estimating depth from stereo pairs and show that it outperforms the state-of-the-art in the high resolution Middlebury imagery [1] as well as in the more challenging KITTI dataset [2].

2 Related Work

In the past few years much progress has been made towards solving the stereo problem, as evidenced by Scharstein et al. overview [21]. Local methods typically aggregate image statistics in a small window, thus imposing smoothness implicitly. Optimization is usually performed using a winner-takes-all strategy, which selects for each pixel the disparity with the smallest value under some distance metric [21]. Traditional local methods [22] often suffer from border bleeding effects or struggle with correspondence ambiguities. Approaches based on adaptive support windows [23, 24] adjust their computations locally to improve performance, especially close to border discontinuities. This results in better performance at the price of more computation.

Hirschmüller proposed semi-global matching [25], an approach which extends polynomial time 1D scanline methods to propagate information along 16 orientations. This reduces streaking artifacts and improves accuracy compared to traditional methods. In this paper we employ this technique to compute a disparity map from which we build our potentials. In [26, 27] disparities are 'grown' from a small set of initial correspondence seeds. Though these methods produce accurate results and can be faster than global approaches, they do not provide dense matching and struggle with textureless and distorted image areas. Approaches to reduce the search space have been investigated for global stereo methods [28, 29] as well as local methods [30].

Dense and accurate matching can be obtained by global methods, which enforce smoothness explicitly by minimizing an MRF-based energy function. These MRFs can be formulated at the pixel level [31], however, the smoothness is then defined very locally. Slatend-plane MRF models for stereo vision were introduced in [4] and have been since very widely used [5, 6, 7, 9, 10, 13]. In the context of this literature, our work has several distinctive features. First, we use a novel model involving "boundary labels", "junction potentials", and "edge ownership". Second, for inference we employ the convex form of the particle norm-product belief propagation [32], which we refer to as particle convex belief propagation (PCBP) [19]. In contrast, some previous works used particle belief propagation (PBP) [33, 34, 10] which correspond to non-convex norm-product with the Bethe entropy approximation. The efficiency and convexity of PCBP makes it possible to evaluate our approach on hundreds of high-resolution images [2], whereas previous empirical evaluations of slanted-plane models have largely been restricted to the low-resolution versions of the small number of highly controlled Middlebury images. Third, we use a training algorithm based on primal-dual approximate inference [35] which allow us to effectively learn the importance of each potential.

3 Continuous MRFs for stereo

In this section we describe our approach to joint reasoning of boundary labels and depth. We reason at the segment level, employing a richer representation than a discrete disparity label. In particular, we formulate

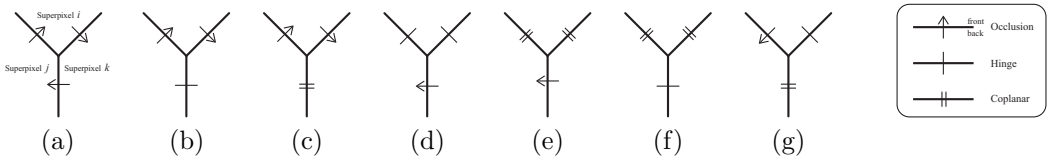


Figure 1: Impossible cases of 3-way junctions. (a) 3 cyclic occlusions, (b) hinge and 2 occlusion with opposite directions, (c) coplanar and 2 occlusion with opposite directions, (d) 2 hinge and occlusion, (e) 2 coplanar and occlusion, (f) 2 coplanar and hinge, (g) hinge, coplanar, and occlusion (superpixel with coplanar boundary is in front).

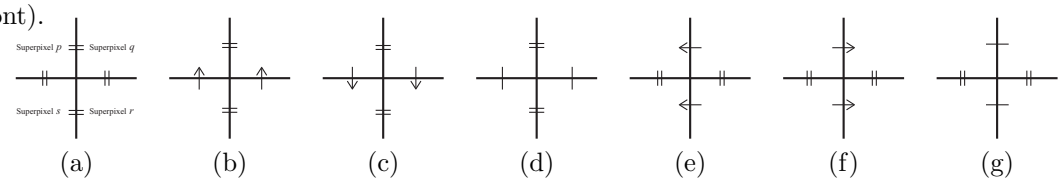


Figure 2: Valid 4-way junctions. (a) 4 coplanar boundaries, (b)-(d) 2 coplanar vertical boundaries and 2 occlusion/hinge horizontal boundaries, (e)-(g) 2 coplanar horizontal boundaries and 2 vertical occlusion/hinge boundariese. A 4-way junction only appears in a region of uniform color.

the problem as inference in a hybrid conditional random field, which contains continuous and discrete random variables. The continuous random variables represent, for each segment, the disparities of all pixels contained in that segment in the form of a 3D slanted plane. The discrete random variables indicate for each pair of neighboring segments, whether they are co-planar, they form a hinge or there is a depth discontinuity (indicating which plane is in front of which).

More formally, let $y_i = (\alpha_i, \beta_i, \gamma_i) \in \Re^3$ be a random variable representing the $i-$th slanted 3D plane. We can compute the disparities of each pixel belonging to the $i-$th segment as follows

$$\hat{d}_i(\mathbf{p}, \mathbf{y}_i) = \alpha_i(u - c_{ix}) + \beta_i(v - c_{iy}) + \gamma_i \tag{1}$$

with $\mathbf{p} = (u, v)$, and $\mathbf{c}_i = (c_{ix}, c_{iy})$ the center of the i-th segment. We have defined γ_i to be the disparity in the segment center as it improves the efficiency of PCBP inference. Let $o_{i,j} \in \{co, hi, lo, ro\}$ be a discrete random variable representing whether two neighboring planes are coplanar, form a hinge or an occlusion boundary. Here, lo implies that plane i occludes plane j, and ro represents that plane j occludes plane i.

We define our hybrid conditional random field as follows

$$p(\mathbf{y}, \mathbf{o}) = \frac{1}{Z} \prod_i \psi_i(\mathbf{y}_i) \prod_\alpha \psi_\alpha(\mathbf{y}_\alpha) \prod_\beta \psi_\beta(\mathbf{o}_\beta) \prod_\gamma \psi_\gamma(\mathbf{y}_\gamma, \mathbf{o}_\gamma)$$

where $\mathbf{y}$ represents the set of all 3D slanted planes, and $\mathbf{o}$ the set of all discrete random variables. The unitary potentials are represented as ψ_i, while $\psi_\alpha, \psi_\beta, \psi_\gamma$ encode potential functions over sets of continuos, discrete or mixture of both types of variables. Note that $\mathbf{y}$ contains three random variables for every segments in the image, and there is a random variable $o_{i,j}$ for each pair of neighboring segments.

In the following, we describe the different potentials we employed for our joint occlusion boundary and depth reasoning. For clarity, we describe the potentials in the log domain, i.e., $\mathbf{w}^T \phi_i = \log(\psi_i)$ (similarly for potentials over cliques). The weights $\mathbf{w}$ will be learned using structure prediction methods.

3.1 Occlusion Boundary and Segmentation Potentials

Our approach takes as input a disparity image computed by any matching algorithm. In particular, in this paper we employ semi-global block matching [25]. Most matching methods return estimated disparity values on a subset of pixels. Let $\mathcal{F}$ be the set of all pixels whose initial disparity has been estimated, and let $\mathcal{D}(\mathbf{p})$ be the disparity of pixel $\mathbf{p} \in \mathcal{F}$. Our model jointly reasons about segmentation in the form of occlusion boundaries as well as depth. We define potentials for each of these tasks individually as well as potentials which link both tasks. We start by defining truncated quadratic potentials, which we will employ in the

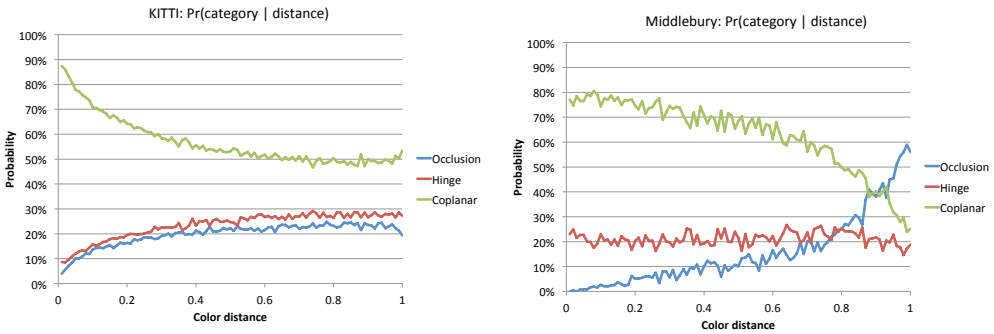


Figure 3: **Color statistics** on (left) KITTI dataset, (right) Middlebury high-resolution



Figure 4: Examples from the KITTI. (Left) Original images. (Right) Disparity errors.

definition of some of our potentials, i.e.,

$$\phi_i^{TP}(\mathbf{p}, \mathbf{y}_i, K) = \min\left(\left|\mathcal{D}(\mathbf{p}) - \hat{d}_i(\mathbf{p}, \mathbf{y}_i)\right|, K\right)^2 \tag{2}$$

with K a constant threshold, and $\hat{d}_i(\mathbf{p}, \mathbf{y}_i)$ the disparity of pixel $\mathbf{p}$ estimated as in Eq. 1. Note that we have made the quadratic potential robust via the min function. We now describe each of the potentials employed in more details.

Disparity potential: We define truncated quadratic unitary potentials for each segment expressing the fact that the plane should agree with the results of the matching algorithm,

$$\phi_i^{\text{seg}}(\mathbf{y}_i) = \sum_{\mathbf{p} \in S_i \cap \mathcal{F}} \phi_i^{TP}(\mathbf{p}, \mathbf{y}_i, K)$$

where S_i is the set of pixels in segment i.

Boundary potential: We employ 3-way potentials linking our discrete and continuous variables. In particular, these potentials express the fact that when two neighboring planes are hinge or coplanar they should agree on the boundary, and when a segment occludes another segment, the boundary should be explained by the occluder. We thus define

$$\phi_{ij}^{\text{bdy1}}(o_{ij}, \mathbf{y}_i, \mathbf{y}_j) = \begin{cases} \sum_{\mathbf{p} \in B_{ij} \cap \mathcal{F}} \phi_i^{TP}(\mathbf{p}, \mathbf{y}_i, K) & \text{if } o_{ij} = lo \\ \sum_{\mathbf{p} \in B_{ij} \cap \mathcal{F}} \phi_j^{TP}(\mathbf{p}, \mathbf{y}_j, K) & \text{if } o_{ij} = ro \\ \frac{1}{2} \sum_{\mathbf{p} \in B_{ij} \cap \mathcal{F}} \phi_i^{TP}(\mathbf{p}, \mathbf{y}_i, K) + \phi_j^{TP}(\mathbf{p}, \mathbf{y}_j, K) & \text{if } o_{ij} = hi \vee co \end{cases}$$

where B_{ij} is the set of pixels around the boundary (within 2 pixels of the boundary) between segments i and j.

	> 2 pixels		> 3 pixels		> 4 pixels		> 5 pixels	
	Non-Occ	**Occ**	**Non-Occ**	**Occ**	**Non-Occ**	**Occ**	**Non-Occ**	**Occ**
GC+occ [36]	41.92 %	43.64 %	34.04 %	35.85 %	29.79 %	31.59 %	27.05 %	28.82 %
OCV-BM [37]	29.18 %	31.65 %	24.94 %	27.32 %	23.21 %	25.48 %	22.02 %	24.18 %
CostFilter [38]	28.46 %	29.96 %	20.74 %	22.13 %	17.24 %	18.47 %	15.40 %	16.50 %
GCS [26]	21.87 %	23.76 %	14.21 %	15.92 %	10.61 %	12.12 %	8.55 %	9.90 %
GCSF [39]	20.75 %	22.69 %	13.02 %	14.77 %	9.48 %	11.02 %	7.48 %	8.84 %
SDM [27]	19.01 %	20.89 %	11.95 %	13.65 %	8.98 %	10.47 %	7.35 %	8.66 %
ELAS [30]	11.25 %	13.71 %	7.60 %	9.77 %	5.91 %	7.79 %	4.87 %	6.52 %
OCV-SGBM [25]	12.48 %	14.86 %	7.40 %	9.54 %	5.51 %	7.43 %	4.52 %	6.21 %
Ours	**9.03** %	**11.58** %	**4.47** %	**6.66** %	**3.13** %	**5.05** %	**2.51** %	**4.23** %

Table 1: Comparison with the state of the art KITTI dataset

	> 1 pixel		> 2 pixels		> 3 pixels		> 4 pixels		> 5 pixels	
	N.-Occ	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**
GC+occ [36]	23.8 %	-	16.6 %	-	13.9 %	-	12.5 %	-	11.5 %	-
EBP [40]	14.3 %	-	10.3 %	-	9.4 %	-	9.0 %	-	8.7 %	-
GCS [26]	13.2 %	-	9.0 %	-	7.4 %	-	6.5 %	-	5.9 %	-
SDM [27]	12.8 %	-	9.3 %	-	8.2 %	-	7.7 %	-	7.3 %	-
ELAS [30]	7.1 %	23.4 %	4.7 %	17.2 %	3.9 %	14.3 %	3.5 %	12.8 %	3.2 %	11.8 %
OCV-SGBM [25]	7.0 %	24.4 %	5.9 %	22.9 %	5.5 %	22.4 %	5.3 %	22.1 %	5.2 %	21.8 %
Ours	**5.1** %	**17.5** %	**3.3** %	**14.3** %	**2.8** %	**13.0** %	**2.5** %	**12.2** %	**2.3** %	**11.6** %

Table 2: Comparison with the state-of-the-art on Middlebury high-resolution imagery. The baselines are provided by the author of [30]

Compatibility potential: We introduce an additional potential which ensures that the discrete occlusion labels match well the disparity observations. We do so by defining $\phi_{ij}^{\text{occ}}(\mathbf{y}_{\text{front}}, \mathbf{y}_{\text{back}})$ to be a penalty term which penalizes occlusion boundaries that are not supported by the data

$$\phi_{ij}^{\text{occ}}(\mathbf{y}_{\text{front}}, \mathbf{y}_{\text{back}}) = \begin{cases} \lambda_{\text{imp}} & \text{if } \exists \mathbf{p} \in B_{ij}: \quad \hat{d}_i(\mathbf{p}, \mathbf{y}_{\text{front}}) < \hat{d}_j(\mathbf{p}, \mathbf{y}_{\text{back}}) \\ 0 & \text{otherwise} \end{cases}$$

We also define $\phi_{ij}^{\text{neg}}(\mathbf{y}_i)$ to be a function which penalizes negative disparities

$$\phi_{ij}^{\text{neg}}(\mathbf{y}_i) = \begin{cases} \lambda_{imp} & \text{if } \min_{\mathbf{p} \in B_{ij}} \hat{d}_i(\mathbf{p}, \mathbf{y}_i) < 0 \\ 0 & \text{otherwise} \end{cases}$$

We impose a regularization on the type of occlusion boundary, where we prefer simpler explanations (i.e., coplanar is preferable than hinge which is more desirable than occlusion). We encode this preference by defining $\lambda_{\text{occ}} > \lambda_{\text{hinge}} > 0$. We thus define our computability potential

$$\phi_{ij}^{\text{bdy2}}(o_{ij}, \mathbf{y}_i, \mathbf{y}_j) = \begin{cases} \lambda_{\text{occ}} + \phi_{ij}^{\text{neg}}(\mathbf{y}_i) + \phi_{ij}^{\text{neg}}(\mathbf{y}_j) + \phi_{ij}^{\text{occ}}(\mathbf{y}_i, \mathbf{y}_j) & \text{if } o_{ij} = lo \\ \lambda_{\text{occ}} + \phi_{ij}^{\text{neg}}(\mathbf{y}_i) + \phi_{ij}^{\text{neg}}(\mathbf{y}_j) + \phi_{ij}^{\text{occ}}(\mathbf{y}_j, \mathbf{y}_i) & \text{if } o_{ij} = ro \\ \lambda_{\text{hinge}} + \phi_{ij}^{\text{neg}}(\mathbf{y}_i) + \phi_{ij}^{\text{neg}}(\mathbf{y}_j) + \frac{1}{|B_{ij}|} \sum_{\mathbf{p} \in B_{ij}} \Delta d_{ij} & \text{if } o_{ij} = hi \\ \phi_{ij}^{\text{neg}}(\mathbf{y}_i) + \phi_{ij}^{\text{neg}}(\mathbf{y}_j) + \frac{1}{|S_i \cup S_j|} \sum_{\mathbf{p} \in S_i \cup S_j} \Delta d_{i,j} & \text{if } o_{ij} = co \end{cases}$$

with $\Delta d_{i,j} = (\hat{d}_i(\mathbf{p}, \mathbf{y}_i) - \hat{d}_j(\mathbf{p}, \mathbf{y}_j))^2$.

Junction Feasibility: Following work on occlusion boundary reasoning [15, 16], we utilize higher order potentials to encode whether a junction of three planes is possible. We refer the reader to Fig. 1 for an

	Super-pixels	> 2 pixels		> 3 pixels		> 4 pixels		> 5 pixels	
		N.-Occ	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**
UCM	91.7	58.06%	58.85%	52.13%	52.86%	48.67%	49.35%	46.26%	46.88%
SLIC	978.6	9.21%	11.73%	4.60%	6.75%	3.26%	5.15%	2.63%	4.32%
SLIC	1198.8	9.16%	11.76%	4.62%	6.86%	3.26%	5.23%	2.62%	4.39%
UCM+SLIC	1191.7	9.03%	11.58%	4.47%	6.66%	3.13%	5.05%	2.51%	4.23%

Table 3: Difference in performance when employing different segmentation methods to compute superpixels on the KITTI dataset. Note that employing an intersection of SLIC+UCM superpixels works best for the same amount of superpixels.

illustration of these cases. We thus define the compatibility of a junction $\{i, j, k\}$ to be

$$\phi^{\text{jct}}_{ijk}(o_{ij}, o_{jk}, o_{ik}) = \begin{cases} \lambda_{\text{imp}} & \text{if impossible case} \\ 0 & \text{otherwise} \end{cases}$$

We also defined a potential encoding the feasibility of a junction of four planes (see Fig. 2) as follows

$$\phi^{crs}_{pqrs}(o_{pq}, o_{qr}, o_{rs}, o_{ps}) = \begin{cases} \lambda_{\text{imp}} & \text{if impossible case} \\ 0 & \text{otherwise} \end{cases}$$

Note that, although these potentials are high order, they only involve variables with small number of states, i.e., 4 states.

Potential for color similarity: Finally, we employ a simple color potential to reason about segmentation, which is defined in terms of the χ-squared distance between color histograms of neighboring segments. This potential encodes the fact that we expect segments which are coplanar to have similar color statistics (i.e., histograms), while the entropy of this distribution is higher when the planes form an occlusion boundary or a hinge. Note that this trend is shown in Fig. 3 (left) for the KITTI [2] dataset. The statistics are less meaningful in the case of the Middelbury high resolution imagery [1], as this dataset is captured in a control environment. We thus reflect these statistics in the following potential

$$\phi^{\text{col}}_{ij}(o_{ij}) = \begin{cases} \min\left(\kappa \cdot \chi^2(h_i, h_j), \lambda_{\text{col}}\right) & \text{if } o_{ij} = co \\ \lambda_{\text{col}} & \text{otherwise} \end{cases}$$

with κ a scalar and $\chi^2(h_i, h_j)$ the χ-squared distance between the color histograms of segments i and j.

3.2 Inference in Continuous MRFs

Now that we have defined the model, we can turn our attention to inference, which is defined as computing the MAP estimate as follows

$$\text{(inference)} \quad \arg\max_{\mathbf{y},\mathbf{o}} \frac{1}{Z} \prod_i \psi_i(\mathbf{y}_i) \prod_\alpha \psi_\alpha(\mathbf{y}_\alpha) \prod_\beta \psi_\beta(\mathbf{o}_\beta) \prod_\gamma \psi_\gamma(\mathbf{y}_\gamma, \mathbf{o}_\gamma) \tag{3}$$

Inference in this model is in general NP hard. Our inference is also particularly challenging since, unlike traditional MRF stereo formulations, we have defined a hybrid MRF, which reasons about continuous as well as discrete variables.

While there is a vast literature on discrete MRF inference, only a few attempts have focussed on solving the continuous case. The exact MAP solution can only be recovered in very restrictive cases. For example when the potentials are quadratic and diagonally dominated, an algorithm called Gaussian Belief propagation

	Super-pixels	> 1 pixel		> 2 pixels		> 3 pixels		> 4 pixels		> 5 pixels	
		N.-Occ	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**
UCM	259.0	35.5%	48.2%	26.7%	37.8%	21.3%	31.4%	17.7%	27.1%	15.2%	23.7%
SLIC	1787.6	5.4%	17.8%	3.6%	14.4%	3.0%	13.1%	2.8%	12.3%	2.6%	11.7%
SLIC	2066.1	5.3%	17.8%	3.5%	14.5%	3.0%	**13.0%**	2.7%	**12.2%**	2.5%	**11.6%**
UCM+SLIC	2042.6	**5.1%**	**17.5%**	**3.3%**	**14.3%**	**2.8%**	**13.0%**	**2.5%**	**12.2%**	**2.3%**	**11.6%**

Table 4: Performance changes when employing different segmentation methods to compute superpixels on the Middlebury high-resolution dataset. Employing an intersection of SLIC+UCM superpixels works best for the same amount of superpixels.

[41] returns the optimal solution. For general potentials, one can approximate the messages using mixture models, or via particles. In this paper we make use of particle convex belief propagation (PCBP) [19], a technique that is guarantee to converge and gradually approach the optimum. This works very well in practice, yielding state-of-the-art results.

PCBP is an iterative algorithm that works as follows: For each random variable, particles are sampled around the current solution. These samples act as labels in a discretized MRF which is solved to convergence using convex belief propagation [32]. The current solution is then updated with the MAP estimate obtained on the discretized MRF. This process is repeated for a fixed number of iterations. In our implementation, we use the distributed message passing algorithm of [42] to solve the discretized MRF at each iteration. Algorithm 1 depicts PCBP for our formulation. At each iteration, to balance the trade off between exploration and exploitation, we decrease the values of the standard deviations $\sigma_\alpha, \sigma_\beta$ and σ_γ of the normal distributions from which the plane random variables are drawn.

Algorithm 1 PCBP for stereo estimation and occlusion boundary reasoning

Set N
Initialize slanted planes $\mathbf{y}_i^0 = (\alpha_i^0, \beta_i^0, \gamma_i^0)$ via local fitting $\forall i$
Initialize $\sigma_\alpha, \sigma_\beta$ and σ_γ
for $t = 1$ to #iters **do**
  Sample N times $\forall i$ from $\alpha_i \sim \mathcal{N}(\alpha_i^{t-1}, \sigma_\alpha)$, $\beta_i \sim \mathcal{N}(\beta_i^{t-1}, \sigma_\beta)$, $\gamma_i \sim \mathcal{N}(\gamma_i^{t-1}, \sigma_\gamma)$
  $(\mathbf{o}^t, \mathbf{y}^t) \leftarrow$ Solve the discretized MRF using convex BP
  Update $\sigma_\alpha^c = \sigma_\beta^c = 0.5 \times \exp(-c/10)$ and $\sigma_\gamma^c = 5.0 \times \exp(-c/10)$
end for
Return $\mathbf{o}^t$, $\mathbf{y}^t$

3.3 Learning in Continuous MRFs

We employ the algorithm of [20] for learning. Given a set of training images and corresponding depth labels, the goal of learning is to estimate the weights which minimize the surrogate partition loss. However, our learning problem, as opposed to the one defined in [20], contains a mixture of continuous and discrete variables. Therefore the surrogate partition loss in our setting requires to integrate over the continuous variables. We note that our continuous variables have robust quadratic potentials, thus integrating over them can be efficiently estimated by discretizing the continuous variables. In practice, summing over 30 particles gives a good approximation for the integral.

4 Experimental Evaluation

We perform exhaustive experiments on two publicly available datasets: Middebury high resolution images [1] as well as the more challenging KITTI dataset [2]. For all experiments, we employ the same parameters which have been validated on the training set. We use a disparity difference threshold $K = 5.0$ pixels, and set $\lambda_{\text{occ}} = 15$, $\lambda_{\text{hinge}} = 3$, $\lambda_{\text{imp}} = 30$ and $\lambda_{\text{col}} = 30$. For the color potential, we use a color histogram with 64 bins

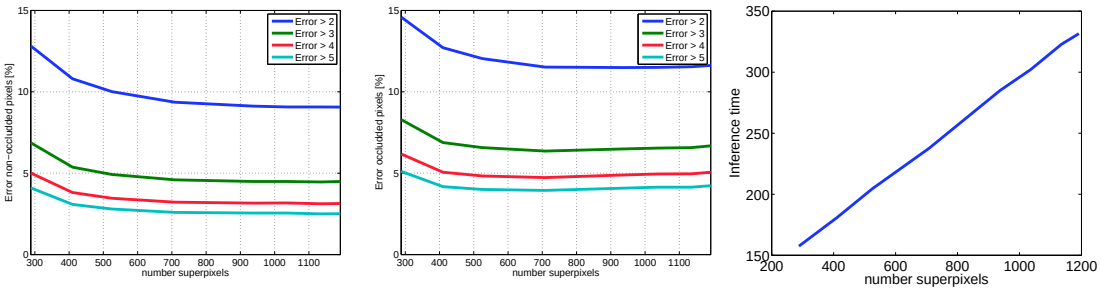


Figure 5: **Importance of the number of superpixels:** KITTI results as a function of the number of superpixels. Even with a small number our approach still outperforms the baselines. (Right) The inference time scales linearly with the number of superpixels

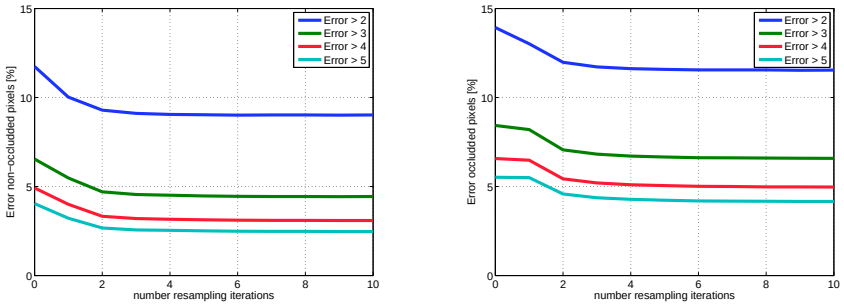


Figure 6: **Importance of the number of re-sampling iterations:** on KITTI.

and set $\kappa = 60$. Unless otherwise stated, we employ 10 particles and 5 iterations of re-sampling for PCBP [19], and run each iteration of convex BP to convergence. For learning, we use a value of C equal to the number of examples and unless otherwise stated use a CRF, i.e., $\epsilon = 1$. We learned the importance of each potential, thus 6 parameters. We employ two different metrics. The first one measures the average number of non-occluded pixels which error is bigger than a fixed threshold. To test the extrapolation capabilities of the different approaches, the second metric computes the average number of pixels (including the occluding ones) which error is bigger than a fixed threshold.

We now describe the characteristics of the databases we evaluate our approach on. Our first dataset consists on high resolution images from the Middebury dataset [1], which have an average resolution of 1239.2×1038.0 pixels. We employ 5 images for training (i.e., *Books, Laundry, Moebius, Reindeer, Bowling2*) and 9 images for testing (i.e., *Cones, Teddy, Art, Aloe, Dolls, Baby3, Cloth3, Lampshade2, Rocks2*). We also evaluate our approach on the KITTI dataset [2], which is the only real-world stereo dataset with accurate ground truth. It is composed of 194 training and 195 test high-resolution images (1237.1×374.1 pixels) captured from an autonomous driving platform driving around in a urban environment. The ground truth is generated by means of a Velodyne sensor which is calibrated with the stereo pair. This results in semi-dense ground truth covering approximately 30 % of the pixels. We employ 10 images for training, and utilize the remaining 184 images for validation purposes.

Comparison with the state-of-the-art: We begin our experimentation by comparing our approach with the state-of-the-art. Table 1 depicts results of our approach and the baselines in terms of the two metrics for the KITTI dataset. Note that our approach significantly outperforms all the baselines in all settings (i.e., thresholds bigger than 2, 3, 4 and 5 pixels). Table 2 depicts similar comparisons for high resolution Middlebury. Once more, our approach outperforms significantly the baselines in all settings. Fig. 4 depicts an illustrative set of example results for the KITTI dataset. Note that despite the challenges that the images

Number of training images	> 2 pixels		> 3 pixels		> 4 pixels		> 5 pixels	
	Non-Occ	**Occ**	**Non-Occ**	**Occ**	**Non-Occ**	**Occ**	**Non-Occ**	**Occ**
1	9.14 %	11.49 %	4.58 %	6.55 %	3.23 %	4.92 %	2.61 %	4.09 %
5	9.02 %	11.58 %	4.46 %	6.66 %	3.13 %	5.06 %	2.50 %	4.24 %
10	9.03 %	11.58 %	4.47 %	6.66 %	3.13 %	5.05 %	2.51 %	4.23 %
20	9.03 %	11.59 %	4.46 %	6.66 %	3.12 %	5.05 %	2.49 %	4.22 %

Table 5: **Training set size:** Estimation errors as a function of the training set size. Note that very few images are needed to learn good parameters.

	> 2 pixels		> 3 pixels		> 4 pixels		> 5 pixels	
	N.-Occ	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**
Oracle	1.38%	1.70%	1.03%	1.27%	0.90%	1.10%	0.82%	0.99%
Initial fit	10.69%	13.72%	6.12%	8.83%	4.59%	7.02%	3.76%	5.98%
Ours	9.64%	12.24%	5.14%	7.35%	3.70%	5.63%	2.97%	4.69%

Table 6: **Oracle performance:** Oracle, our approach and initial fit on KITTI.

pose, our approach does a good job at estimating disparities.

Segmentation strategy: We next investigate how the segmentation strategy affects the stereo estimation. Towards this goal we evaluate the results of our approach when employing UCM segments [43], SLIC superpixels [44] or the intersection of both as input. Table 3 depicts results on the KITTI dataset. UCM performs very poorly as the number of superpixels on average is very small, and some of the superpixels are very large. Therefore, a single 3D plane is a poor representation for the disparities in those large segments. SLIC performs quite well, but the intersection of SLIC and UCM superpixels outperforms the other strategies. This is also expected, as UCM respects the boundaries much better than SLIC. Note that as shown in Table 4 similar results are observed for the Middlebury dataset.

Number of superpixels: We next investigate how well our approach scales with the number of superpixels in terms of computatinal complexity as well as accuracy. Fig. 5 shows results for the KITTI dataset when varying the number of superpixels. Our approach reduces performance gracefully when reducing the amount of superpixels. Note that inference scales linearly with the number of superpixels, taking on average 5.5 minutes per high resolution image when employing 1200 superpixels and 2.5 minutes when using 300.

Number of re-sampling iterations: We evaluate the effects of varying the number of resampling iterations on the performance of our approach. As shown in Fig. 6, our approach converges to a good local optima after only 2 resampling iterations. This reduces the inference cost from 5.5 minutes per high-resolution image for 5 iterations to 2.2 minutes for 2 iterations.

Training set size: We evaluate the effect of increasing the training set size in Table 5. Even when training with a single image we outperform all baselines.

Oracle performance: We next evaluate the best performance that our model can achieve, by fitting the model to the ground truth disparities. This shows an upper-bound on the performance that our method could ever achieve if we were able to learn an energy that has its MAP at the ground truth, and if we were able to solve the NP-hard inference problem. Tables 6 and 7 depict the oracle performance in terms of both the occluded an non-occluded pixels for both datasets. Note that as KITTI does not release the test ground truth, we compute this values using 10 images for training and the rest of the training set for testing. We also report performance of our initialization which is computed by fitting a local plane to the results

	> 1 pixel		> 2 pixels		> 3 pixels		> 4 pixels		> 5 pixels	
	N.-Occ	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**	**N.-Occ**	**Occ**
Oracle	2.0%	6.2%	1.4%	5.5%	1.3%	5.3%	1.2%	5.2%	1.1%	5.1%
Initial fit	6.1%	21.8%	4.2%	19.3%	3.5%	18.3%	3.2%	17.7%	3.0%	17.3%
Ours	5.1%	17.5%	3.3%	14.3%	2.8%	13.0%	2.5%	12.2%	2.3%	11.6%

Table 7: **Oracle performance:** Oracle, our approach and initial fit on Middlebury.

Noise	RMS (pixels)	boundary error
0	0.44	0.3 %
1	0.80	0.6 %
2	1.37	1.9 %
3	2.24	5.3 %
5	4.40	8.9 %

Table 8: **Robustness to noise**: RMS as well as boundary error as a function of noise.

of semi-global block matching [25]. Note that the oracle can achieve great performance, showing that the errors due to the 3D slanted plane discretization are negligible.

Robustness to noise: We investigate the robustness of our approach to noise by building a synthetic dataset, which is composed of 10 images for training and 90 images for test. Each Image has a resolution of 320×240 pixels and contains several planes. The average number of superpixels is 108.0. We create $\mathcal{D}(\mathbf{p})$ by sampling 3 to 5 points at random on the boundaries and generating disparities by corrupting the ground truth with Gaussian noise of varying standard deviation. Table 8 shows RMS errors for disparity as well as percentage of boundary variables wrongly estimated.

Family of structure prediction problems: We evaluate the performance of our learning algorithm as a function of its parameter ϵ which ranges from CRFs for $\epsilon = 1$ to structural SVMs for $\epsilon = 0$. Fig. 9 depicts performance on the KITTI dataset. Our approach results in state-of-the-art performance for all settings.

Task Loss: We evaluate the importance of incorporating different tasks loss in the learning framework of [20]. In particular, we employ RMS as well as the same loss that we employ for evaluation. Note that the loss has little effect.

5 Conclusion

We have presented a novel stereo slanted-plane MRF model that reasons jointly about occlusion boundaries as well as depth. We have formulated the problem as inference in a hybrid MRF composed of both continuous (i.e., slanted 3D planes) and discrete (i.e., occlusion boundaries) random variables, which we have tackled using particle convex belief propagation. We have demonstrated the effectiveness of our approach on high resolution imagery from Middlebury as well as the more challenging KITTI dataset. In the future we plan to investigate alternative inference algorithms as well as other segmentation potentials.

References

[1] Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. IJCV **47** (2002) Microsoft Research Technical Report MSR-TR-2001-81, November 2001.

[2] Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? In: CVPR. (2012)

	> 2 pixels		> 3 pixels		> 4 pixels		> 5 pixels	
	Non-Occ	**Occ**	**Non-Occ**	**Occ**	**Non-Occ**	**Occ**	**Non-Occ**	**Occ**
$\epsilon = 0.0$	9.16 %	11.84 %	4.64 %	6.96 %	3.25 %	5.30 %	2.58 %	4.42 %
$\epsilon = 0.5$	9.15 %	11.59 %	4.59 %	6.65 %	3.24 %	5.02 %	2.59 %	4.17 %
$\epsilon = 1.0$	9.03 %	11.58 %	4.47 %	6.66 %	3.13 %	5.05 %	2.51 %	4.23 %

Table 9: **Family of structure prediction problems:** for the KITTI dataset. Note that $\epsilon = 0$ is structural SVM and $\epsilon = 1$ is CRF.

	> 2 pixels		> 3 pixels		> 4 pixels		> 5 pixels	
	Non-Occ	**Occ**	**Non-Occ**	**Occ**	**Non-Occ**	**Occ**	**Non-Occ**	**Occ**
ℓ_2 loss	9.03 %	11.58 %	4.47 %	6.66 %	3.13 %	5.05 %	2.51 %	4.23 %
>2 pixels	9.03 %	11.59 %	4.46 %	6.65 %	3.13 %	5.05 %	2.50 %	4.24 %
>3 pixels	9.05 %	11.58 %	4.47 %	6.64 %	3.13 %	5.03 %	2.49 %	4.20 %
>4 pixels	9.02 %	11.58 %	4.47 %	6.67 %	3.13 %	5.06 %	2.50 &	4.23 %
>5 pixels	9.05 %	11.60 %	4.48 %	6.66 %	3.13 %	5.04 %	2.50 %	4.22 %

Table 10: **Loss functions**: Performance changes when employing different loss functions for learning in the KITTI dataset

[3] Tappen, M., Freeman, W.: Comparison of graph cuts with belief propagation for stereo, using identical mrf parameters. In: ICCV. (2003)

[4] Birchfield, S., Tomasi, C.: Multiway cut for stereo and motion with slanted surfaces. In: ICCV. (1999)

[5] Hong, L., Chen, G.: Segment-based stereo matching using graph cuts. In: CVPR. (2004)

[6] Bleyer, M., Gelautz, M.: A layered stereo matching algorithm using image segmentation and global visibility constraints. ISPRS Journal of Photogrammetry and Remote Sensing **59** (2005) 128150

[7] Klaus, A., Sormann, M., Karner, K.: Segment-based stereo matching using belief propagation and a self-adapting dissimilarity measure. In: ICPR-06. (2006)

[8] Deng, Y., Yang, Q., Lin, X., Tang, X.: A symmetric patch-based correspondence model for occlusion handling. In: ICCV. (2005)

[9] Yang, Q., Engels, C., Akbarzadeh, A.: Near real-time stereo for weakly-textured scenes. In: British Machine Vision Conference. (2008)

[10] Trinh, H., McAllester, D.: Unsupervised learning of stereo vision with monocular cues. In: Proc. BMVC, Citeseer (2009)

[11] Zitnick, C., Kanade, T.: A cooperative algorithm for stereo matching and occlusion detection. PAMI **22** (2000)

[12] Kolmogorov, V., Zabih, R.: Multi-camera scene reconstruction via graph cuts. In: ECCV. (2002)

[13] Bleyer, M., Rother, C., Kohli, P.: Surface stereo with soft segmentation. In: CVPR. (2010)

[14] Waltz, D.L.: Generating semantic description from drawings of scenes with shadows. MIT Artificial Intelligence Laboratory (1972) Workin Paper 47.

[15] Malik, J.: Interpreting line drawings of curved objects. International Journal of Computer Vision **1** (1987) 73–103

[16] Hoiem, D., Stein, A.N., Efros, A.A., Hebert, M.: Recovering occlusion boundaries from a single image. In: ICCV. (2007)

[17] Ashutosh Saxena, J.S., Ng, A.Y.: Depth estimation using monocular and stereo cues. In: IJCAI. (2007)

[18] Maire, M., Arbelaez, P., Fowlkes, C., Malik, J.: Using contours to detect and localize junctions in natural images. In: CVPR. (2008)

[19] Peng, J., Hazan, T., McAllester, D., Urtasun, R.: Convex max-product algorithms for continuous mrfs with applications to protein folding. In: International Conference on Machine Learning (ICML). (2011)

[20] Hazan, T., Urtasun, R.: A primal-dual message-passing algorithm for approximated large scale structured prediction. Advances in Neural Information Processing Systems **23** (2010) 838–846

[21] Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. IJCV (2002)

[22] Konolige, K.: Small vision system. hardware and implementation. In: International Symposium on Robotics Research. (1997) 111–116

[23] Kanade, T., Okutomi, M.: A stereo matching algorithm with an adaptive window: Theory and experiment. ICRA (1994)

[24] Yoon, K.j., Member, S., Kweon, I.S.: Adaptive support-weight approach for correspondence search. PAMI **28** (2006) 650–656

[25] Hirschmueller, H.: Stereo processing by semiglobal matching and mutual information. PAMI **30** (2008) 328–41

[26] Cech, J., Sara, R.: Efficient sampling of disparity space for fast and accurate matching. In: BenCOS. (2007)

[27] Kostkova, J.: Stratified dense matching for stereopsis in complex scenes. In: BMVC. (2003)

[28] Wang, L., Jin, H., Yang, R.: Search space reduction for mrf stereo. In: ECCV. (2008)

[29] Veksler, O.: Reducing search space for stereo correspondence with graph cuts. In: BMVC. (2006)

[30] Geiger, A., Roser, M., Urtasun, R.: Efficient large-scale stereo matching. Computer Vision–ACCV 2010 (2010) 25–38

[31] Sun, J., Zheng, N.N., Shum, H.Y.: Stereo matching using belief propagation. PAMI **25** (2003)

[32] Hazan, T., Shashua, A.: Norm-product belief propagation: Primal-dual message-passing for approximate inference. Information Theory, IEEE Transactions on **56** (2010) 6294–6316

[33] Koller, D., Lerner, U., Angelov, D.: A general algorithm for approximate inference and its application to hybrid bayes nets. In: UAI. (1999)

[34] Ihler, A., McAllester, D.: Particle belief propagation. In: AISTATS-09. (2009)

[35] Hazan, T., Urtasun, R.: A primal-dual message-passing algorithm for approximated large scale structured prediction. In: NIPS. (2010)

[36] Kolmogorov, V., Zabih, R.: Computing visual correspondence with occlusions using graph cuts. In: ICCV. (2001) 508–515

[37] Bradski, G.: The opencv library. Dr. Dobb's Journal of Software Tools (2000)

[38] Rhemann, C., Hosni, A., Bleyer, M., Rother, C., Gelautz, M.: Fast cost-volume filtering for visual correspondence and beyond. In: CVPR. (2011)

[39] Cech, J., Sanchez-Riera, J., Horaud, R.P.: Scene flow estimation by growing correspondence seeds. In: CVPR. (2011)

[40] Felzenszwalb, P.F., Huttenlocher, D.P.: Efficient belief propagation for early vision. IJCV **70** (2006)

[41] Weiss, Y., Freeman, W.: Correctness of belief propagation in gaussian graphical models of arbitrary topology. Neural Computation **13** (2001)

[42] Schwing, A., Hazan, T., Pollefeys, M., Urtasun, R.: Distributed message passing for large scale graphical models. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2011)

[43] P.Arbelaez, M.Maire, C.Fowlkes, J.Malik: Contour detection and hierarchical image segmentation. In: PAMI. (2011)

[44] Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S.: Slic superpixels. Technical Report 149300 EPFL (2010)